Regular Paper

Toward a Systematic Approach for Selection of NASA Technology Portfolios

Charles R. Weisbin,* Guillermo Rodriguez, Alberto Elfes, and Jeffrey H. Smith

Jet Propulsion Laboratory, California Institute of Technology, 4800 *Oak Grove Drive, Pasadena,* CA 91109

Received 8 December 2003; Accepted 24 May 2004, after one or more revisions
Published online in Wiley InterScience (www.interscience.wiley.com).
DOI 10.1002/sys.20010

ABSTRACT

There is an important need for a consistent analytical foundation supporting the selection and monitoring of R&D tasks that support new system concepts that enable future NASA missions. This capability should be applicable at various degrees of abstraction, depending upon whether one is interested in formulation, development, or operations. It should also be applicable to a single project, a program comprised of a group of projects, an enterprise typically including multiple programs, and the overall agency itself. Emphasis here is on technology selection and new initiatives, but the same approach can be generalized to other applications, dealing, for example, with new system architectures, risk reduction, and task allocation among humans and machines. The purpose of this paper is to describe one such approach, which is in its early stages of implementation within NASA programs, and to discuss several illustrative examples. © 2004 Wiley Periodicals, Inc. Syst Eng 7: 285–302, 2004

Key words: technologies; prioritization; missions; return on investment

1. INTRODUCTION

The current drive toward "One NASA," a goal that embodies cross-enterprise Agency Missions and an Integrated Space Plan, has created an important need for an overall integrated agency-wide approach to systems analysis. A central element of such an approach is the development of a consistent methodological foundation for selecting and monitoring R&D tasks that support new system concepts to enable or enhance future missions.

This capability should be applicable at various degrees of abstraction, depending upon whether one is interested in formulation, development, or operations. It should also be applicable to a single project, a pro-

*Author to whom all correspondence should be addressed (e-mail: charles.r.weisbin@jpl.nasa.gov).

Contract grant sponsor: National Aeronautics and Space Administration.

gram comprised of a group of projects, an enterprise typically including multiple programs, and NASA itself.

START (STrategic Assessment of Risk and Technology) offers one approach [Chase et al., 2003; Dolgin and Weisbin, 2003; Elfes et al., 2003; Lincoln et al., 2003; Smith, Wertz, and Weisbin, 2003a: 101; Smith, Wertz, and Weisbin, 2003b; Smith, Dolgin, and Weisbin, 2003; Neff et al., 2004] that is in its early stages toward achieving this capability. Developed within the Strategic Systems Technology Program Office, a division of the Office of the Chief Technologist at Jet Propulsion Laboratory, START offers systems for quantifying the features of each development candidate, assessing its risk, and calculating its probable return-on-investment. The approach is currently under further development and evaluation at various programmatic and institutional organization levels within NASA.

Emphasis here is on technology selection and analysis of technological options that are in their early stages of development, but a similar approach is envisioned to deal with such other issues of selection of new system architectures and risk reduction during system integration during the entire life cycle of a mission design process.

2. METHODOLOGY

The following describes the general procedure that the START team follows. It represents a significant departure from the process by which many important decisions about funding and technology selection have been made until now.

Though expert decision-makers may be guided by extensive experience and good judgment, they have human limitations. Usually, a decision-maker will consider only a few attributes when comparing competing technologies. Our system's usefulness, as much as anything, is that it induces decision-makers to consider all of the pertinent attributes, and provides a sound method for using them in the decision-making process.

Even when a decision-maker is confident about a selection based solely on his or her experience and judgment, the START process can provide a valuable, objective foundation to support that decision.

Please note, however, that not all studies begin at Step 1 and continue through to Step 8. A sponsor may have determined the answers to early-stage questions before initiating a study. Or a study may focus on, for example, identifying and evaluating possible system architectures for a given mission (Step 3).

In some cases, we may be called upon to assess the usefulness of a particular technology that was funded as basic research. To take a hypothetical example, the developer of a particular nanotechnology might want to know how it could be put to use in NASA's various programs. In such a case, we would employ a "bottom-up" approach, beginning at Step 5 and working upward to Step 2.

Frequently, we are called upon to split the difference: working top-down until we have derived the capability requirements for a particular mission, then switching to bottom-up to identify the capabilities of a particular set of technologies that were funded as basic research. The case study, "Rover Autonomy #2," described below, is a good example of this approach. The action lies in matching capabilities with capability requirements.

2.1. Develop a Clear, Complete Statement of the Problem To Be Studied

State the problem unambiguously, specifying what is to be maximized or minimized, with all pertinent policy, schedule, and budget constraints. Probe to uncover any unstated assumptions that need to be taken into account. Unarticulated assumptions can undermine a study.

In a study of competing technologies, for example, the decision-maker can specify options about top-level policy. One policy could be to fund only as many technologies as can be brought to completion. Another policy might prefer to fund all of the competing technologies at some level. There can also be a weighted combination of multiple objectives such as an appropriate weighting of the two specific policies presented above. This preference about policy can guide the subsequent technology selection analysis, by providing a framework within which the analysis is conducted.

Most often, we are asked to address the problem of maximizing science return subject to a given resource. However, our studies are capable of pursuing any number of other objectives, such as minimizing cost for variable range in performance, maximizing continuity of tasks, maximizing public interest, etc.

2.2. Identify Top-Level Goal

Identify top-level goals and quantify what would constitute satisfying those goals. For example, a mission to detect possible life 1 km below the Martian surface would be one way to meet NASA's goal of searching for life on other worlds. For NASA work, we draw goals, investigations, and experiments from NASA strategic plans and science working group meeting reports.

2.3. Develop or Select One or More Architectures for Accomplishing the Goal

Design or select architectures, including precise scenarios, for conducting specific subsets of the desired experiments. A study may address mission architectures, system architectures, or both. For example, for the goal described above, a mission architecture might include launching a spacecraft, landing it safely in a certain location on Mars, having a rover disembark and travel to where scientists suspect a pool of underground water, drilling to a depth of 1 km, retrieving a sample, analyzing the sample for signs of life, and reporting the results to Earth. A system architecture may be limited to the design and functions of the rover.

The START team can also help sponsors identify the time horizon they wish to target for development of their technologies. For example, estimated mission science return can be based on projected Code S and Code Y missions as depicted in their respective roadmaps from 2009 through 2025.

2.4. Identify the Capabilities Needed for the Architecture

Decompose the mission or system concepts into specific quantitative capability requirements whose importance is based on their estimated contribution to the objective stated in Step 1 (such as maximizing science return). Our models are capable of capturing interdependencies between capabilities. For example, a Mars rover's sample acquisition capability depends on coordination of its sensing and manipulation capabilities.

It should be noted that it is common within NASA to have science working groups prioritize mission goals and measurements to achieve those goals, and these prioritizations should be reflected in the objective functions. However, no claim is made that all scientists agree with these lists; science return continues to be largely defined in subjective terms, and these lists are updated frequently as new information becomes available and the scientific focus of a mission changes. More generally, the objective function can incorporate other concerns of the decision-maker, such as overall cost, human risk, mission longevity, etc. Selection of the terms of the objective function and the associated utilities or weights is the responsibility of the decision-maker, since they have to reflect his preferences.

2.5. Identify Technologies That Could Provide the Needed Capabilities

Assess technology candidates that purport to fulfill or partially fulfill the required capabilities. Capture uncertainties in their capabilities, using performance attributes and their probability distributions. Define each technology development task by at least four critical metrics:

a. Performance requirement attributes
b. Budget estimate
c. Scheduled delivery date
d. Risk level

2.6. Evaluate and Rank the Technology Candidates To Identify Which To Use or Fund for Development

Rank technologies by calculating their contributions to all relevant capabilities and missions. Generate uniform unitless values to compare attributes with dissimilar metrics (for example, mass in kg, volume in cm^3, cost in dollars, etc.).

Risk may be calculated and considered, both in terms of an individual technology's risk of failure (useful in comparison with competing technologies), and in terms of the impact a technology's failure would have on the entire mission. The risk is not only due to a given technology's failure, but it is also due to the technology's possible inability to satisfy all the desired goals.

Construct optimal portfolios (sets of technologies for the desired purpose) for the objective stated in Step 1 (such as maximizing the total science return within allowable cost limits and other programmatic constraints).

2.7. Validate Results

Though it is impossible to compare a study's outcome with "truth," we consider our results validated if they are consistent with all known information (experiments, models, expert opinion, uncertainties). If not, we reexamine the inputs and model assumptions that led to the study's result.

2.8. Track and Reconstitute the Technology Portfolio as Needed

Maintain an optimal portfolio as technologies mature and customer requirements change.

3. VALIDATION

Validation is the process of comparing a model's output with a real system or, lacking one, with an expert's judgment. If the result is consistent with all known information and the expert's opinion, we consider it validated. A positive validation confirms that the model's output represents the most reasonable result, within the limits of uncertainty.

If a result is invalidated, we examine the model's assumptions, revisit the inputs to see whether they were estimated accurately, and/or adjust for any new constraints that were not previously expressed.

It is important to note that some degree of uncertainty surrounds every input, and some inputs, such as the relative importance of a particular attribute, can only be assessed by experts, and are likely to have relatively high variability. For technologies that do not yet exist, virtually all inputs may have to be estimated amid considerable uncertainty.

If the experts involved in a study's validation process reaffirm the values for each attribute, the decision-maker may reconsider a conflicting opinion and bring it into accord with the study's results. Alternatively, the experts and decision-maker may revise some of the input values, leading to a different outcome.

One of the inevitable limitations in attempting to validate our methodology is that the actual outcome of decisions about portfolio selection is difficult if not impossible to quantify. System engineers typically define validation to mean that the thing that is validated is tested, and the test results are compared with the requirements the thing is intended to satisfy. The intent of such validation is to compare and test predictions with what actually occurs. Although desirable, it is not feasible to validate a portfolio based on what is observed in the future. Specific component results can be observed to validate pieces of data. However, it is simply not feasible to compare predictions with actual data, until a process such as the one we describe has been utilized in an operational setting long enough to allow sufficient data about impact of portfolio selection decisions to be gathered.

3.1. Sensitivity and Uncertainty

We can calculate which attributes (such as mass, volume, cost, or an aspect of performance) were most influential in producing the study's outcome, versus some other particular outcome. In practice, this is most useful when a study's outcome differs from the outcome preferred or expected by a decision-maker.

If a small change in the value assigned a particular attribute would produce a large difference in the result, that attribute is said to have high sensitivity. Conversely, low sensitivity indicates that even a big change in the value assigned a given attribute would have little impact on the study's results.

Relative uncertainty in a result is deduced from the product of sensitivity and uncertainty in the data that led to the result. If an attribute's uncertainty is much higher than that of the other attributes, it may be worthwhile to try to reduce that level of uncertainty. If all attributes have about the same level of uncertainty, we focus on sensitivity.

We can use sensitivity information in two ways. First, with the dominant influences on the study's output brought to light, a decision-maker can decide whether these particular influences make sense. If, for example, the cost of testing has high sensitivity in a study of competing technologies, but the decision-maker does not think that the cost of testing should be much of a determining factor, that is a signal that we need to reconsider the factors that produced such a high sensitivity for that attribute.

Sensitivity results allow us to calibrate the model by identifying attributes to target for reevaluation of the input values. Minor revisions to a few highly sensitive attributes values may make the results more accurate. Once the model is calibrated, the decision-maker can be confident of the results, even ones which contradict prior intuition.

The goal of this process, however, is not simply to make the study agree with an expert's preconceived ideas. It is to examine the underlying reasons for the difference in outcomes, and to determine whether any of the initial values should be changed on their own merits.

This procedure exposes the implications and ramifications of any given result, whether it is the study's initial output or the expert's preference. Result "A" means that all the values, preferences, and weightings that led to "A" are the best choices. Result "B" means that all the parameters that led to an output of "B" are the best choices. Going through this process leads a decision maker to examine those values, preferences, and weightings, and to make sure that they are as accurate as they can be.

In doing so, we build a solid foundation for whatever result the study ultimately produces. If, after this reexamination process, the study confirms the decision-maker's original preference, it provides a comprehensive explanation for why that is the best prediction that can be made. On the other hand, if it leads to a change of mind, the decision-maker will know exactly why such a change was warranted.

4. CASE STUDIES

Following is a group of case studies that involve technology tasks applied to exploration of the surface of Mars. The techniques employed, however, are applicable to a wide variety of mission types inside and outside of NASA.

The common thread in these studies is the development of models that enable us to calculate the impact technologies would have on the science return of their missions. This enables us to assign values to the projected return-on-investment for each technology, a very useful tool in ranking the technologies for funding and development.

Our studies therefore are intended for sponsors who are interested in examining the various technological options that may be available at an early stage for broad classes of missions. In response to this, we focused on identifying individual technologies, and groupings of technologies, whose integration has not occurred as yet, but that could occur at a later stage in a mission design and implementation process. While focusing on the identification of advanced technology options for a range of missions, we recognize that other risks besides technology development risk may drive the design process in later stages of a mission design evolution, including systems integration, testing, and even ultimately in risk inherent in mission operations. These integrated system risks, and not only technological risk, have been the primary sources of recent failures in operational flight systems ranging shuttle, to space telescopes, and planetary surface exploration vehicles. Case studies that cover this broader set of risks, including those associated with technologies that need inflight validation are discussed in [Neff et al., 2004].

4.1. Autonomy for Mars Rovers

Whenever current rover systems experience a failure, they stop, wait for the next scheduled opportunity to communicate its problem to Earth (relatively brief periods each day, due to limitations of the rover's solar batteries), and wait for new commands attempting to resolve the problem. After each command, the Earthbound controllers await Pathfinder's progress report before issuing a follow-up command.

Technology that would increase a rover's autonomy—that is, improve its ability to conduct science while reducing its need to phone home for help—would save a great deal of time and therefore enable the rover to accomplish much more.

Following are two case studies [Elfes et al., 2003; Lincoln et al., 2003] that represent efforts to determine the relative benefits of investing in various software technologies that purport to help Mars rovers do science more efficiently, avoid most failures, and diagnose and correct their own problems when failures occur.

The first study (Rover Autonomy #1) focuses on technologies that were proposed specifically to reduce fault rates observed during extensive field-testing in Mars-like terrain here on Earth.

The second study (Rover Autonomy #2) analyzes technologies that were funded as basic research, only loosely coupled to a mission. Hence, we needed to determine technology-derived capabilities and match those capabilities with mission requirements. These technologies are more advanced than those studied in Rover Autonomy #1, capable of automating entire sequential operations.

4.2. Case Study 1: Rover Autonomy #1

We conducted this study [Elfes et al., 2003] to determine the relative benefits of developing various autonomy software technologies for a surface rover in the proposed Mars Science Laboratory (MSL) mission scheduled for 2009. Since the rover prototypes had been extensively field-tested in Mars-like terrain on Earth, we had access to an extensive body of real-world information.

We decomposed the mission into functional steps (acquire panorama, develop range map, plan path, etc.) covering long-range traverse, short-range approach to target, and sample acquisition and handling. For each of these steps in each mission element, we noted the kinds and frequencies of failure, and the time that was lost while the controllers developed a strategy to mitigate the failure.

For each of the science operations (moving samples to the rover's onboard analytic lab, conducting contact experiments, moving to a new site, etc.), we developed a utility function that corresponds to the effectiveness in contributing to the mission goals [Lincoln et al., 2003]. The utility function was endorsed by a science group, and reflects the preferences of the decision-makers. For example, the first sample collected in a bag may be worth 40% of the total mission value. In general, intrusive experiments, such as grinding up a rock sample and analyzing it with a mass spectrometer, merited the highest values.

Technology development cost data was provided by the technology developers and underwent independent review [Lincoln et al., 2003].

We calculated the abilities of the autonomy software technologies under study to mitigate potential failures, as well as the difficulty in developing responses to each of them. Subsequent work transformed the difficulty estimation into dollars. Since the cost of each technology cannot be predicted with certainty, we established uncertainty estimates in return-on-investment with regard to performance and, through modeling, to science return.

Each autonomy software technology was judged by two attributes: ability to save time (measured in Martian days, or "sols"), and cost. The relative contributions of the autonomy technologies appear in Figure 1.

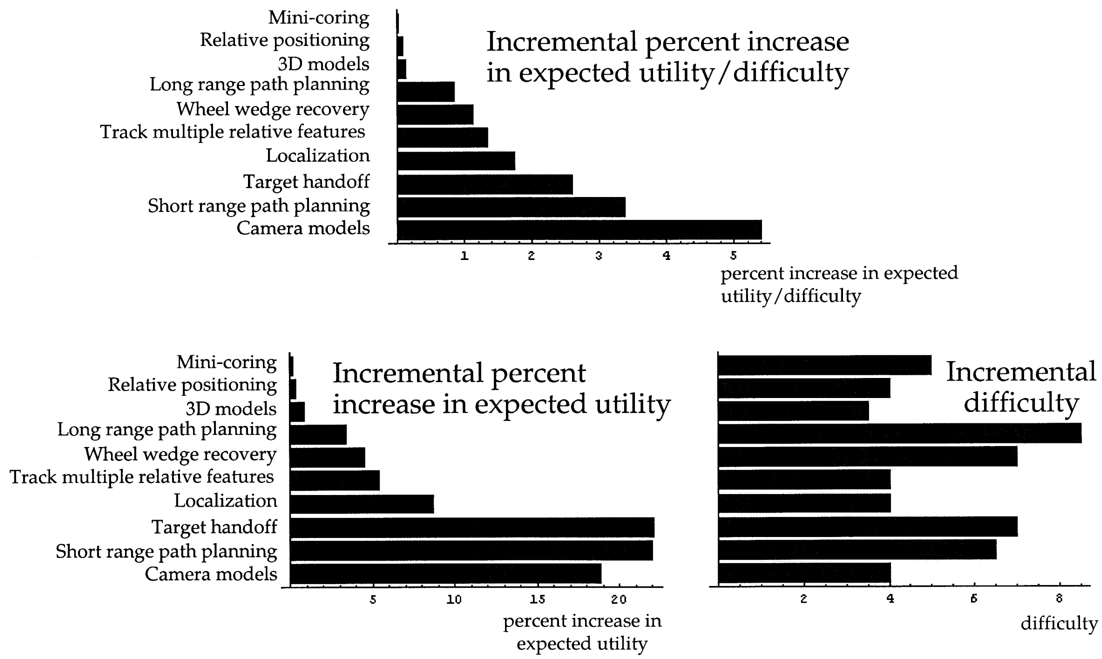

Figure 1. Relative contribution of autonomy technologies. The lower left chart shows our prediction of the relative utility of autonomy software for each technology, while the chart at lower right shows the projected difficulty in developing it (which relates to cost). Dividing utility by difficulty (or benefit by cost) yields the return-on-investment data illustrated in the top chart.

4.3. Case Study 2: Rover Autonomy #2

This pilot study analyzed two potential missions to Mars, one to an equatorial region and a subsequent one to a north polar region. These two mission concepts are consistent with current advanced Mars mission plans over the next two decades.

The study's objective was to develop and demonstrate a system for evaluating and comparing groups of advanced autonomy software technologies, most of which are currently at the R&D stage (TRL 2 or 3), in terms of the impact they would have on science return in each of the two missions.

We began with a top-down methodology in which we derived, from top-level mission goals, the technology capabilities that would be required to operate at the two disparate sites and produce the desired science. Then we went bottom-up, taking a group of technologies, determining their capabilities, and evaluating the potential impact that enhanced autonomy would enable them to make.

4.3.1. Deriving Capability Requirements

Following the top-down methodology, we decomposed the mission goals into their constituent functional requirements. For example, the science requirements call for such capabilities as mobility, instrument placement, sample acquisition, and telecommunications. Mobility, to take one of these capabilities, entails range-mapping to estimate distance to nearby objects and possible hazards, path planning, and obstacle avoidance. Obstacle avoidance, in turn, breaks down into obstacle detection and navigation. Using this system, we ultimately arrived at a comprehensive list of functional requirements that could benefit from enhanced autonomy provided by the technologies under consideration.

4.3.2. Autonomy Technologies Selected for Evaluation

Turning to the "bottom-up" portion of our methodology, we selected 15 technology groups as representatives of the diverse technological interests, and of the seven technology areas that enable surface operations (see Table I).

The task before us was to match their capabilities with the capability requirements derived from the mission goals, and to determine the extent to which each of these technology groups would produce greater science results if it had enhanced autonomy.

4.3.3. Calculating Impact on Science Return

Through interviews with experts, we developed performance parameters for each of the technology groups,

Table I. CICT Technology Group

Area	Technology Group
1 and 2	Onboard Fault Identification for Planetary Rovers
1 and 2	System-Level Verification Technology
1 and 2	Autonomy Infusion Simulation Environment
3	Distributed Control Testbed for Autonomy
3	Rover Autonomy Architecture
4	Single-Cycle Instrument Placement
4	Rover-Based Manipulation
5	Multimedia Human Computer Interfaces
5	Human-Centered Computing for MER
6	Onboard Science Analysis
7	MER Rover Sequence Generation
7	Contingency Planning for Concurrent Activities
7	Accelerated Long-Range Traverse
7	System for Mobility and Access to Rough Terrain
7	Super-Resolved 3-D Surface Models from Rover Images

Key to Areas:

1. Fault Management
2. Validation/Verification
3. Software Architecture
4. Approach/Instrument Placement
5. Human-Computer Interaction
6. Sample Handling
7. Mobility

and determined which elements of the mission the technology would help.

For example, we determined the impact that "System for Mobility and Access to Rough Terrain" would have on the rover's traverse rate. Then we plugged that information into the Mission Model (see Fig. 2) to calculate its impact on the number of sols (Martian days) that this technology would save over the current state-of-the-art, as a percentage of the total mission duration.

Saving sols means enabling the rover to spend time doing science instead of—in this case—traveling. So saved sols are presumed to correlate to an increase in science value. We used the number of saved sols in the ROI Model (see Fig. 3) to determine each technology's return on investment, and the resulting numbers were used to rank the technologies.

Table II shows the results of the initial prioritization. The task names have been replaced by the letters A–O because the data is still preliminary and under review. We provide the results table, with the data and supporting models, to all parties involved to begin a dialogue on the perceived impact and rationale.

ROI represents increase in science value (as measured by the number of sols saved over state-of-the-art, or "SOA") divided by cost. When calculating the *combined* ROI for each technology task, we gave the MSL value twice as much weight as the polar value. This weighting is somewhat arbitrary and could be changed if desired. But it was intended to reflect the fact that

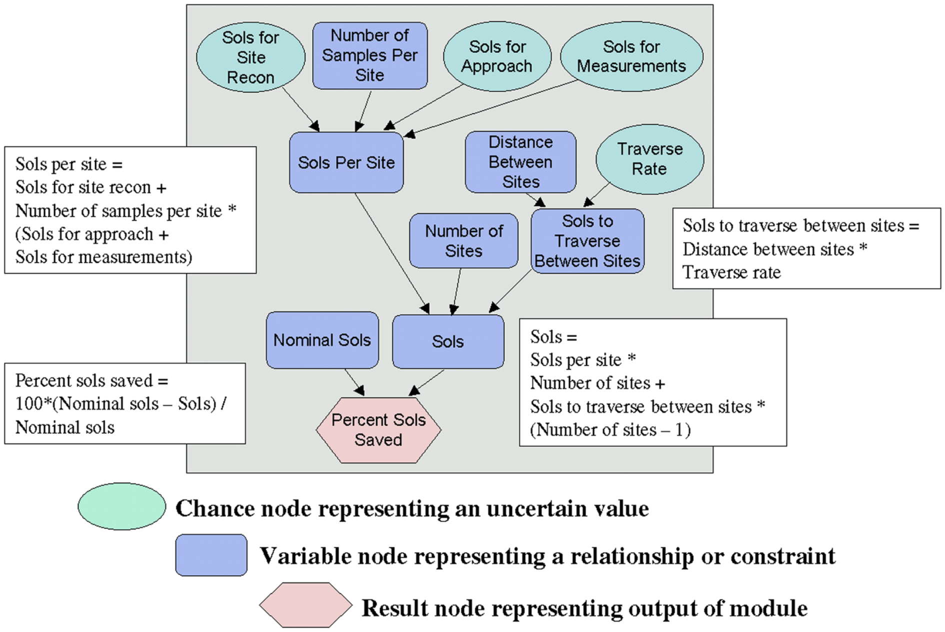

Figure 2. Mission Model. This template illustrates the procedure for calculating how many sols each technology would save, as a percentage of the total mission. Each technology would impact (and presumably improve) one or more of the bubbles that lead to determining a number of "sols." "Nominal sols" refers to the number of sols that would be spent using state-of-the-art technology as represented by the MER rovers. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

these technologies are more likely to be used in the more imminent MSL mission, and to be precursors to the technologies that will enable and enhance the polar mission. Though these technologies are innovative, far exceed SOA for the most part, and are intended for long-term impact, they will have as much as a decade for further improvement between the two missions.

Note also that these ROI numbers are not intended to represent final, definitive evaluations, but rather a solid basis for further investigation and discussion.

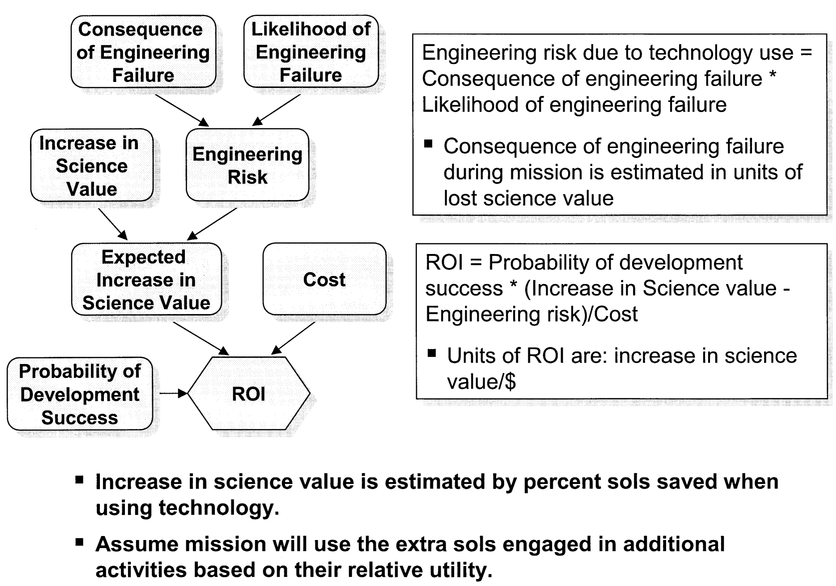

Figure 3. Return on Investment (ROI) model. This template illustrates the procedure for determining ROI for each technology group. "Increase in science value" uses the "percent sols saved" number Calculated in the Mission Model.

Table II. Initial Results

Task	MSL ROI	Polar ROI	Combined ROI
Technology 1	4	6	5
Technology 2	7	9	8
Technology 3	15	17	15
Technology 4	14	17	15
Technology 5	10	11	10
Technology 6	13	10	12
Technology 7	31	32	31
Technology 8	3	3	3
Technology 9	4	8	5
Technology 10	4	0	3
Technology 11	23	46	31
Technology 12	7	15	10
Technology 13	14	14	14
Technology 14	5	23	11
Technology 15	2	2	2

Combined ROI used illustrative weighting with relative ratio 2:1 for MSL and Polar missions.

They indicate the potential performance of each technology under certain conditions and for specific purposes. A given technology might benefit additional operations that, if factored into the study, would improve the technology's ROI. Similarly, we could amplify the study by factoring in additional metrics—such as development and operations cost, heritage value, innovation, and public inspiration—and potentially arrive at different results.

However, the study does demonstrate that it is possible to estimate mission-level science return impacts of diverse autonomy technologies, that the results can be very useful in assisting decision-makers in the selection of technology groups for funding and development, and that these methods are applicable to a wider class of technologies and mission classes.

4.4. Case Study 3: Predicting the Cost of New Technologies

Investigators seeking technology development funding typically cast the best light on their estimates of how much time and money they will require. Add to this trait the fact that technologists and mission designers often have conflicting, unexpressed assumptions about what is required, and you have the makings of costly misunderstandings and cost overruns.

This case study concept [Smith, Wertz, and Weisbin, 2003b] was aimed at developing a process to generate plausible cost estimates grounded on clear assumptions.

We developed a process (see Fig. 4) for estimating the cost of new technology that included uncertainty and an independent peer review of the estimate. It is based on interviews with technology representatives that focus on cost and performance relationships for each technology:

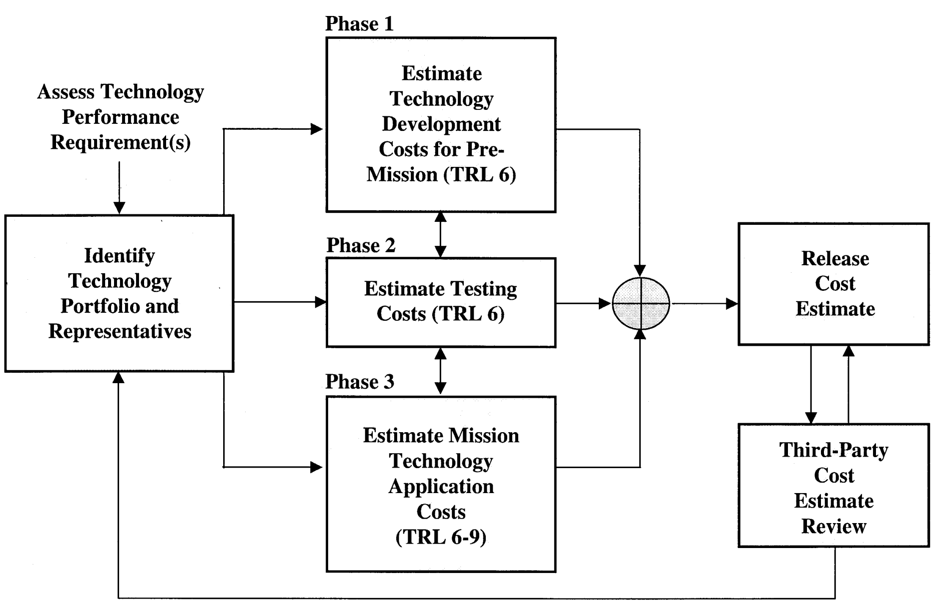

Figure 4. Process to estimate R&D costs. This particular task subset dealt only with technologies up to TR 6, and so did not include the action described in the lower middle box.

1. What are the important relationships that influence the cost?
2. What are the development issues?
3. What happens to performance if the cost is higher or lower?
4. What happens to cost if performance is higher or lower?
5. What assumptions underlie the cost estimate?
6. What is the probability of successfully developing the technology?

As a test case, we applied the process to a set of autonomy software technologies for Mars rovers that were the focus of the "Rover Autonomy #1" study.

The interviews in this case revealed important and subtle factors such as technology interdependencies, resource dependencies, and areas of common problems for the technologies studied. The third-party review was critical in helping to (1) validate the original prediction, (2) identify missing or redundant cost issues affecting the initial prediction, and (3) determine any adjustments that might need to be made to the original cost estimate.

Figure 5 illustrates the likelihood of success of three of the tasks as a function of Available Resources. While the task was to model the relationships between performance, cost, and schedule for autonomy software, the general approach should be extensible to other technologies, including hardware systems.

An optimal $50 million portfolio does not necessarily simply add new technologies to those of a $40 million portfolio. Expanding the budget may make an entirely different set of technologies possible and preferable.

By more reliably predicting the costs of component technologies and considering the inter-relationships of their science return, we can help decision-makers to determine the best place to set the cutoff points for their technology budgets. Together, Figures 5 and 6 can help a decision maker to optimize a portfolio.

Suppose he or she has about $2 million to spend on autonomy software technology. Considering the three technologies represented on these graphs, the decision-maker can fund one of three possible portfolios:

1. Camera models and target handoff. But there will only be enough money to fund target handoff to the point where the top graph indicates less than a 0.4 probability of being completed.
2. Target handoff alone, but to the level where the top graph indicates near certainty that it will be completed.
3. Short range path planning, but only to the level where it has around a 0.5 probability of being completed.

Figure 6 tells us that Portfolio #1 will save about 15 sols for the camera models plus about 10 sols for the target handoff, for a total of 25 sols saved. Portfolio #2 would save about 35 sols. Portfolio #3 would save about 11 sols.

All other things being equal, the best return-on-investment would come from portfolio #2, which would save 35 sols with a near-certainty of completion.

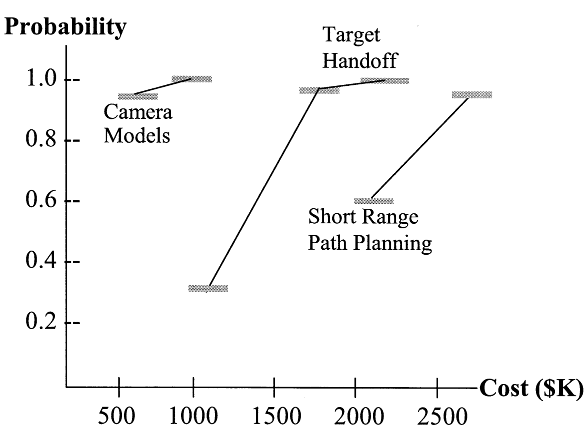

Figure 5. Likelihood of R&D success as a function of available resources. The graph to the left shows the probability of completing three tasks to their specified level of performance, at a range of budgets. For example, the probability of completing target handoff rises from about 0.3 at roughly $1.1 million to about 0.95 at a cost of about $1.75 million. The green shading around each budget point indicates the amount of uncertainty in the figure.

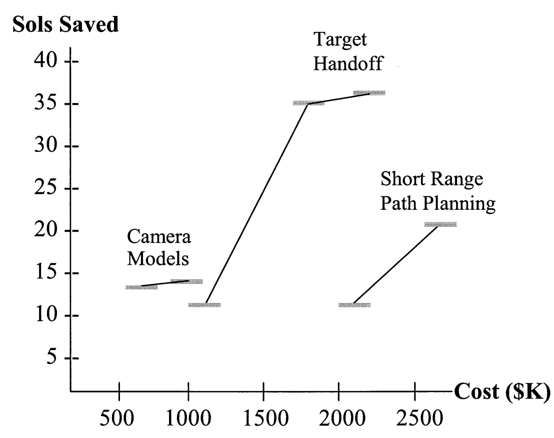

Figure 6. Effective science return as a function of R&D investment. This graph to the right indicates the performance level (measured in the number of Martian days, or sols, that would be saved) one would expect at the budget levels plotted in the previous graph. For the target handoff, the number of sols saved increases from about 10 at roughly $1.1 million to about 35 at about $1.75 million. The data for both graphs were derived from interviews with experts.

4.5. Case Study 4: Optimizing Technology Portfolios for Mars Missions

This study [Smith, Dolgin, and Weisbin, 2003] illustrates a more extensive approach to developing optimal technology portfolios for specific budgets. Mars program goals include discovering whether life ever arose there, determining the planet's climate history and the evolution of its surface and interior, and preparing for human missions. We began our study by developing concepts for missions to accomplish these goals during the timeframe of 2009–2020. They are summarized in Table III.

Next, we developed quantitative capability requirements to enable the potential missions, and identified the technology development efforts required to enable those capabilities, taking note of their funding levels, probabilities of success, and the alternate technologies available for use if the new technology cannot be successfully developed. Technology costs and mission costs were included to screen infeasible mission sets and technology portfolios from the solution.

We picked three levels of technology investment for a 12-year period—$25 million per year, $50 million per year, and $75 million per year—and used an optimization program to determine which sets of technology would yield the best science return at each funding level. The results appear in the Table IV.

At the $25M/yr technology budget, only 24 out of the 511 portfolios met the budget constraints. One lander/rover and one orbiter had the lowest technology costs that fit within the budget profile. The striking result was that although the emphasis was on landed missions, only one landed mission option was feasible.

At the $50M/yr technology budget, the number of affordable technology portfolios increased to 288 out of 511 possibilities and it allowed 15 additional technologies to enter the solution that enabled two additional missions. From these results it was clear that the $50M/yr budget had opened the tradeoff space between technologies and enabled a variety of missions (*in situ*, sample return, and global orbiters).

All 511 technology portfolios fit within the technology budget constraint at the $75M/yr level, enabling one additional mission. The mission budget constraint coupled with the higher costs and risks due to dependencies on the Mars Science Laboratory prevented the Polar rover and Wildcat missions from entering the optimal solution. It was also found that if the strategy were revised to emphasize landed missions rather than sample return, then landed missions push the sample return mission out of the solution.

The total technology investment costs and mission costs for each of the technology budget levels clearly showed the sensitivity of the optimizing mission set on technology investment.

Technology Budget	$25M/yr	$50M/yr	$75M/yr
Technology Cost	$73M	$214M	$238M
Mission (Program) Cost	$1430M	$3580M	$4070M
Number of Missions Possible	2	4	5

4.6. Case Study 5: Lander vs. Rover

This case study [Elfes et al., 2003] compares the impact of investments in precision landing and long-range roving technologies on a hypothetical mission to Mars. We show how to develop an optimal investment strategy

Table III. Candidate Missions for the Mars Program

Mission Name	Description
Mars Science Laboratory	Mission to measure science measurement in-situ with a rover
Volcanology Rover	Rover mission to characterize volcanic region with in-situ sampling
Synthetic Aperture Radar Orbiter	Orbiter sounding for surface science experiments and mapping
Imaging/Atmospheric Sounding Orbiter	Next generation remote sensing orbiter (imaging and atmospheric sounding)
G. Marconi Orbiter	Telecommunications orbiter relay for high data rate communications
Telesat Orbiter	Small Mars telecommunications orbiter for high data rate communications
MSR Sample Lander	Sample return with a Mars ascent vehicle
Wildcat Lander	Lander with 30mm depth drilling system

Table IV. Technology Resources, Associated R&D, and Mars Missions Enabled

Technology Investment	Technology Portfolio	Missions Enabled
$25M Per Year	• Sample characterization • Rover mobility at 160-200 m • Orbit science resolution • Telcom network, Mars-to-Earth	• Volcanology Rover • Imaging/Atmospheric Sounding orbiter **Mission Set Cost = $1430M** **Technology Portfolio Cost = $73M**
$50M Per Year	• Precision landing • Impact attenuation • Hazard avoidance • Forward planetary protection time – Forward planetary protection, number of organisms – Back planetary protection – Mars orbit rendezvous – Mars ascent vehicle • Sample characterization • Mobility at 230-450 m • Sample handling, contamination • Multi-mission survivability • Approach/Instrument placement • Mobility at 160-200 m • On-orbit science resolution • Telecon network, Mars-to-Earth	• Mars Science Laboratory • Volcanology Rover • Imaging/Atmospheric Sounding orbiter • Mars Sample Return **Mission Set Cost = $3580M** **Technology Portfolio Cost = $214M**
$75M Per Year	All technologies for $50M case plus on-orbit science, wavelength.	• Mars Science Laboratory • Volcanology Rover • Synthetic Aperture Radar orbiter • Imaging/Atmospheric Sounding orbiter • Mars Sample Return **Mission Set Cost = $4070M** **Technology Portfolio Cost = $238M**

that minimizes mission risk, given a fixed total technology investment budget. More details about this study are contained in Smith, Wertz, and Weisbin [2003a].

The baseline mission scenario for this study is a Mars 2009-class mission with precision landing capability and a long-range rover. There are three preselected science sites, including the target-landing site, with a total traversal distance of 6 km. Total mission time is 90 sols (Martian days), with 50 sols allocated to traversal.

The results are shown below in Figure 7, which illustrates a tradeoff between investment in landing and investment in roving technology development. In this graph, investment in lander technology is shown on the horizontal axis, and investment in rover technology is shown on the vertical axis. The dollar amounts on the two axes are connected by diagonal "isobudget" lines. Every point along the straight line that connects $40M on the lander axis with $40M on the rover axis, for example, indicates a combined investment of $40M.

The curved lines represent levels of risk of mission failure. The top curved line, for instance, represents a 10% chance that the mission will fail (or, to put it more optimistically, a 90% probability of success).

The uppermost "risk" curve that is intersected by any given "budget" line indicates the lowest risk level that budget can buy. The point of intersection reveals what combination of investments in lander and rover technology will achieve that lowest possible risk.

For example, if you have $40M to spend, you look along the $40M diagonal line until you see where it intersects the highest risk curve. $40M does not intersect the very top curve, which indicates a 10% risk of failure, but it does intersect the 20% curve. So the least amount of risk you can have for a $40M budget is 20%. And by seeing where that intersection point falls on the two axes, you can determine how that $40M budget

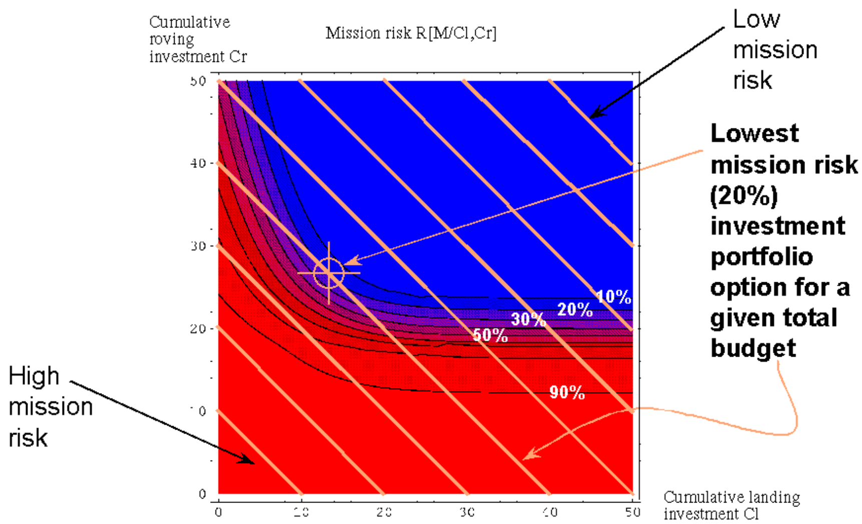

Figure 7. Landing-roving trade space of risk, cost, and performance.

should be divided between lander and rover technology. In this case, you would spend about \$13M on lander technology and about \$27M on rover technology to achieve the lowest possible risk for that budget: 20%.

If risk level is more important to you than dollar amount, you can use this graph to see how much you have to spend—and where you should spend it—to achieve that level of risk. For example, if nothing greater than a 10% risk (that is, nothing less than a 90% probability of success) is acceptable, you can see that the least amount you can budget is about \$44M, of which about \$26M should be spent on rover technology, and about \$18M should be spent on lander technology.

Another method of visualizing the results from this study is shown in Figure 8. On this graph, total budget levels vary vertically. The minimum mission risk achievable at each budget level is shown on the left, while the corresponding technology portfolio appears on the right.

5. RELATIONSHIP TO OTHER WORK

Risk assessment, tracking, and mitigation are themes of core importance not only for the development of space missions, but for other areas as varied as large-scale engineering projects, urban planning, health manage-

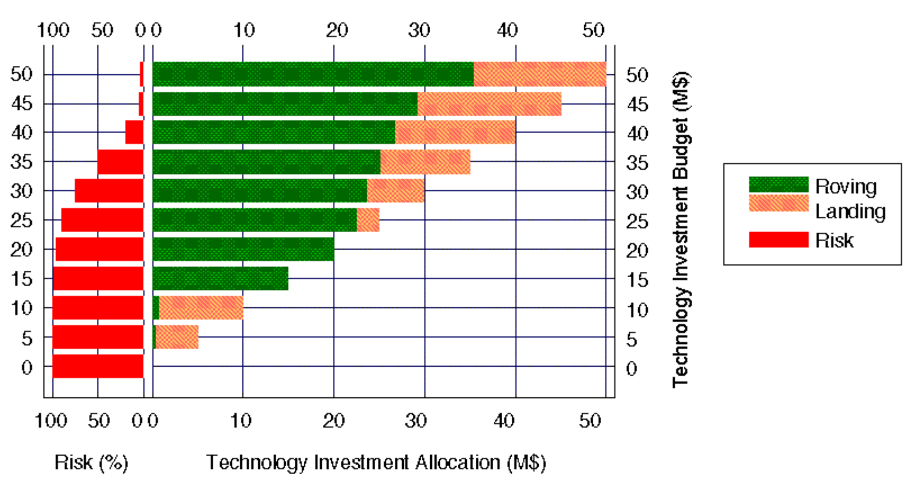

Figure 8. Investment strategies and associated levels.

ment, economic decisions, or military planning. Consequently, risk assessment and mitigation technologies have been developed and applied extensively in the engineering disciplines [Haimes, Kaplan, and Lambert, 2001; Haimes, 1998], flight planning [FAA System Handbook, 2000], economics, medicine, governmental planning, military decision-making [Dombroski et al., 2002: 19; McNichols and Owens, 1995; Steves, 1995; Risk Management Concepts & Guidance, 1989; Intelligence Community R&D Database; Braunstein and Salsamendi, 1994], etc.

While some of the methodologies developed in other areas have been incorporated and applied by NASA, space mission planning, development, launch, and operations present unique challenges for risk assessment and management. These include:

- Most deep space exploration missions are one-of-a-kind, and even "twin" missions such as the MER rovers are faced with significantly different challenges in their respective landing sites. Consequently, the knowledge and experience gained with one mission can only be partially transferred and applied to a new mission.
- Complete system validation is only partially possible before launch and deployment, since deep-space temperature, pressure, acceleration, and radiation conditions cannot in general be fully simulated on Earth.
- Mission lifecycles, going from initial mission concepts to detailed analysis, mission architecture selection, technology development, system integration and testing, and spacecraft deployment and operation, require extensive periods of time. This, as well as the large investments required for any space mission, precludes the ability of doing extensive and statistically meaningful system validation through actual repeated deployments.
- Finally, uncertainties in mission scenarios, schedule and budget delays, long lead times required for mission planning, development and deployment, and mission requirements that often tax technologies to their limit, all contribute very high levels of uncertainty to the whole risk assessment and management process.

While the paper focuses on a set of applications to specific case studies, there are certain features in the underlying methodology in portfolio selection developed over the last several years, and these features allow us to cast the contributions of this paper within the context of the work of others in similar areas of research. Our methodology for Mars portfolio optimization [Smith, Wertz, and Weisbin, 2003a] uses a Monte Carlo simulation method adapted for conditional dependencies to model uncertainty, together with a mathematical programming search algorithm to do constrained optimization over technology portfolios embedded within mission portfolios. This methodology expands the work of others [Faulkner, 1996; Hertzfeld and Vonortas, 1996; Hertzfeld, 1992; Granot and Zuckerman, 1991: 2; Smith and Feinberg, 1994: 218; Czajkowsky and Jones, 1986: 17] that address the same problem, but it offers the following unique features:

- Simulates uncertain development outcomes for technologies with multiple descope options
- Allows dependencies between technologies (if parent fails, dependent technology fails)
- Solves two-stage portfolio problem by (1) evaluating technology portfolios to identify highest value technology for enabling largest number of projects and (2) simultaneously evaluates project portfolios to identify set of projects with highest expected return
- Enumerates all portfolio combinations in search space.

To our knowledge ours may be the first application of such methods to study the optimization of space technology portfolios for Mars missions. Collecting the performance database necessary to apply the advanced methodology and obtain relevant quantitative results constitutes one of the central contributions of this paper.

Similar remarks can be made about the case study describing the trade between landing and roving technologies. This case study draws from the rich body of work [Jensen, 2001; Knight, Glaessgen, and Sleight, 2002; Kotz, Lumelskii, and Pensky, 2003; Pearl, 1988] in decision analysis, with specific emphasis on the use of Bayesian networks and probabilistic reasoning. Our main contribution here is in reducing to engineering practice these generally powerful and broadly applicable methods theoretical methods, using detailed computer simulations [Jain et al., 2003] that accurately predict the interaction of roving and landing vehicles with the physical characteristics of the Mars terrain. It is this combination of advanced methods in probabilistic inference, with detailed physics-based models to estimate the underlying probabilities, that represents the main contribution of this case study. A more detailed description of the technical foundations of this study can be found in [Elfes et al., 2003].

6. SCALE-UP OF TECHNOLOGY RESOURCE ALLOCATION PROCESS

We have illustrated in Table V an overall methodology for systematic technology resource allocation, which

Table V. Activity Description and Schedule for Large Scale Portfolio Analysis Implementation

Activity	Output	Schedule
Objective Function of the Stakeholder(s)	Quantities to be maximized/ minimized and associated constraints	Few Days
Science Value (*NASA Strategic Priorities*)	Weighted set of science objectives and measurements	Two Weeks
Compilation of Advanced Mission Concepts (*Future mission concepts that may not have point design at this stage*)	Advanced mission concepts which might achieve some of the desired measurements 1-2 page top level summaries of selected set of representative mission concepts	One month
Mission Decomposition	Hierarchical decomposition from mission goals to engineering capability requirements	Six Months (Assume 40 mission concepts drawn from existing and current efforts, and teams of 6 per concept)
Requirements Compilation & Generation	Specification of capability requirements, associated uncertainties, schedule, and mission concept cost	Three Months (Assume 40 mission concepts and teams of 6 per concept)
Capability Prioritization	Capabilities can be prioritized at this Stage with approximate costs and risks.	Two Months
Technology Characterization and Forecasts	Projected performance, costs, risks, schedule, dependencies, and uncertainties	Four Months (Assume 600 technologists spending two weeks each, and review teams of 5 independent folks to validate)
Capability/Technology Matching (*WHAT/HOW*)	Technology composition up to capabilities achieved and matching to those desired	Three Months
Portfolio Optimization	Multi-attribute combinatorial optimization	Six Months

we believe can be scaled to the Enterprise and Agency level. We have several suggestions that, if institutionalized within the NASA framework, can facilitate the systematic technology resource allocation process.

1. This type of analysis must be supportive of the decision-makers objective function; i.e., what to maximize/minimize and what to set as constraints. Relative value between programs and objectives needs to be understood; presently such value is frequently established within a given program through interactions of an Advisory Committee with Program management. However, interprogram value considerations are currently dealt with at the agency level, often implicitly rather than explicitly. Designation of value and constraints doesn't require a lot of effort or time, but does represent a cultural change (the information can be kept discrete if necessary).
2. Advanced mission concepts and decomposition of capability requirements (and associated uncertainties) is currently done by each individual project, but such information is not generally available. Such information is essential to the technology resource allocation process, and should be contained within a NASA database that includes advanced mission concept descriptions and objectives.
3. NASA currently has a database (Technology Inventory) intended to characterize the status of technology development. This database can be the underpinnings of our analyses if it is augmented with information from the technologist as follows:

a. The performance objective of each task needs to be quantified with associated metrics and uncertainty bounds.
b. The schedule and estimated cost to complete along with associated uncertainty needs to be provided.
c. Dependencies of one technology upon another, where applicable, should be explicitly noted.
e. A quantified estimate of the probability of task success needs to be provided in order to understand the likelihood of achieving the proposed output, assuming that the task is fully funded; one would not generally expect this to be unity since these are R&D tasks.

Much of this information currently exists, but it is found only in the heads of specific individuals and as such is not easily accessible for comprehensive analysis. Systematically providing such information would greatly facilitate the efficiency of the implementation of the technology resource process for NASA described herein.

This effort is now under review for application to NASA-wide technology allocation decisions at organizational levels responsible for agency-wide technology development and financial management, as well as strategic formulation of an integrated overall NASA architecture vision.

It should be noted that the technology portfolio assessments for the case studies discussed above have concentrated on relevant technologies that are currently being funded by various NASA sources. Although the methodology presented here can directly incorporate information about technologies being developed at other governmental agencies, academia, and industry (see, for example, the Intelligence Community Research and Development Data Base [ICRD]), this was not the focus of the customer decision-makers, who were concerned about making optimal technology portfolio choices within their funding programs.

7. CONCLUSIONS

We have proposed a flexible system that assists decision-makers in evaluating all pertinent attributes of development candidates, including risk and uncertainty, and identifying the main drivers of a result. The system provides a sound foundation for the decision-making process, based on the candidates' predicted contribution to science return or other goals. We have shown that it is quite possible to estimate mission-level science return impacts of diverse technologies, even when those technologies were conceived primarily as basic research.

We have demonstrated a system for making plausible predictions of the cost of new technologies, of determining when diminishing returns make further development inadvisable, and of optimizing technology portfolios at various budget levels.

The case studies cited here illustrate our methodology and the results it can produce. We have emphasized, however, that a study's outcome is generally not intended to be a definitive conclusion, but rather a basis for further investigation and discussion. Ultimately, the process provides solid support for a decision-maker's judgment. The case studies provide a foundation from which more extended applications to programs and missions can be investigated. Such an investigation will be quite challenging and interesting.

8. WEB SITE

The START web site offers many more case studies, and describes how our methodology is applied to other areas besides Mars. Please visit http://start1.jpl.nasa.gov.

ACKNOWLEDGMENT

The research described in this paper was performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

REFERENCES

D.M. Braunstein and M.S. Salsamendi, R&D planning at ARCO Chemical, Res Technol Management (September–October) 37(5) (1994), 33–37.

J.P. Chase, A. Elfes, et al., Identifying technology investments for future space missions, AIAA Space 2003 Conf, Long Beach, CA, September 2003.

A. Czajkowsky and S. Jones, Selecting interrelated R&D projects in space technology planning, IEEE Trans Eng Management 33 (1986), 17–23.

B. Dolgin and C. Weisbin, Building a pathway to Mars: Mars Technology Program analysis and case study, presented at the 2003 Portland Int Conf Management Eng Technol (PICMET): Technol Management Reshaping World, Portland, OR, July 20–24, 2003.

M.H. Dombroski, Y.Y. Haimes, J.H. Lambert, K. Schlussel, and M. Sulcoski, Risk-based methodology for support of operations other than war, Mil Oper Res Sci 7(1) (2002), 19–38.

A. Elfes, W.P. Lincoln, G. Rodriguez, C. Weisbin, and J.A. Wertz, A risk-based framework for optimization of technology portfolio investment choices for early space mission design, IEEE Trans Aerospace Eng (2003), submitted.

FAA System Safety Handbook. Federal Aviation Administration, Washington, DC, 2000.

T.W. Faulkner, Option pricing theory and strategic thinking: Literature review and discussion, Eastman-Kodak Working Paper, May 1996.

D. Granot and D. Zuckerman, Optimal sequencing and resource allocation in R&D projects, Management Sci 37(2) (1991), 140–156.

Y.Y. Haimes, S. Kaplan, and J.H. Lambert, Risk filtering, ranking, and management framework, Risk Anal, to appear.

Y.Y. Haimes, Risk modeling, assessment, and management, Wiley, New York, 1998.

H.R. Hertzfeld and N.S. Vonortas, R&D project selection in the public sector, Space Policy Institute Working Paper, George Washington University, Washington, DC, February 21, 1996.

H.R. Hertzfeld, "Measuring returns to space research and development," Space economics, J.S. Greenberg and H.R.Herzfeld (Editors), AIAA, Washington, DC, 1992.

Intelligence Community Research and Development Data Base, http://ntidb1.gsfc.nasa. gov/icradd.

A. Jain, J. Guineau, C. Lim, W.P. Lincoln, M. Pomerantz, G. Sohl, and R. Steele, Roams: Planetary surface rover simulation environment, presented at the Int Symp Artif Intell, Robotics Automation Space (i-SAIRAS 2003), Nara, Japan, May 19–23, 2003.

F.V. Jensen, Bayesian networks and decision graphs, Springer-Verlag, New York, 2001.

N.F. Knight, E.H. Glaessgen, and D.W. Sleight, An approach to risk-based design incorporating damage-tolerance analysis, AIAA Paper No. 2002-1267, April 2002, pp. 525–543.

S. Kotz, Y. Lumelskii, and M. Pensky. The stress-strength model and its generalizations: Theory and applications, Imperial College Press, London, 2003.

W. Lincoln, A. Elfes, G. Rodriguez, and C.R. Weisbin, Relative benefits of potential autonomy technology investments, Int Conf Space Mission Challenges Inf Technol (SMC-IT), Pasadena, CA, July 13–16, 2003, pp. 47–52.

G.R. McNichols and R.C. Owens, Air Force risk analysis handbook, TR-9007/54-1R, Management Consulting & Research, Inc., January 1995, McLean, VA.

J. Neff, J. Witz, R. Some, and C. R. Weisbin, XML-based tools for assessing potential impact of technology for space validation, presented at IEEE Aerospace Conf, Big Sky, MT, March 2004.

J. Pearl, Probabilistic reasoning in intelligent systems: Networks of plausible inference, Morgan Kaufmann, San Francisco, 1988.

Risk Management Concepts and Guidance. Defense Systems Management College (DSMC), 1989, Reading, MA.

J.H. Smith, J. Wertz, and C. Weisbin, Building a pathway to Mars: Technology investment for science return, J Space Mission Arch (JSMA), Issue 3 (Fall 2003a), pp. 101–115.

J.H. Smith, J. Wertz, and C.R. Weisbin, Predicting the cost of new technology, an approach and case study for autonomy technologies, engineering for complex systems, JPL Internal Document D-26750, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, October 2003b.

J.H. Smith, B. Dolgin, and C.R. Weisbin, Reaching Mars: Multi-criteria R&D portfolio selection for Mars exploration technology planning, Proc 32nd Annu Meeting Western Decision Sci Inst, Marriott Resort and Beach Club, Lihue, Hawaii, April 15–19, 2003, pp. 685–687.

J.H. Smith and A. Feinberg, Reaching for the right portfolio strategy: Multi-criteria R&D planning problems and high-performance computing, Proc 23rd Annu Meeting Western Decision Sci Inst, Maui, Hawaii, March 29–April 2, 1994, pp. 218–219.

M.R. Steves, An examination of Department of Defense risk management policy for nondevelopmental items acquisition programs, March 1995, Naval Postgraduate School, Monterey, CA, March 1995.

Charles R. Weisbin, Eng. Sc. D., 1969 Columbia University. He currently serves as Deputy Program Manager for the Strategic Systems Technology Office of the Chief Technologist at the Jet Propulsion Laboratory. Prior to this, at JPL, he led the Surface Systems Thrust Area, within the NASA Cross-Enterprise Technology Development Program. He was Program Manager for Robotics and Mars Exploration Technology Programs and, before this, Section Manager for Robotic Systems and Advanced Computer Technology. At the Oak Ridge National Laboratory, he was Director of the Robotics and Intelligent Systems Program and Director for the Center for Engineering Systems Advanced Research. Dr. Weisbin was Associate Professor of Computer Science, teaching artificial intelligence at the University of Tennessee. He was the cochairman of the NASA Telerobotics Intercenter Working Group for seven years and received the 1993 NASA Exceptional Service Medal for formulation and development of the NASA Telerobotics Program. He is a recipient of the 1998 Thomas O. Paine Award for Advancement of Human Exploration to Mars and an IEEE Computer Society Golden Core Member award. He also received the 1999 NASA Exceptional Service Medal for development and infusion of robotics from basic research into NASA flight missions and missions of other agencies. Finally, he received the award for Outstanding Leadership in Surface Robotics for the year 2000, the Decadal Planning Team achievement award in 2001, and the NASA Group Achievement Award in 2001 for contributions to the NASA Decadal Planning Team. (charles.r.weisbin@jpl.nasa.gov)

Guillermo Rodriguez is a research engineer at the Jet Propulsion Laboratory (JPL) of the California Institute of Technology. He received a Ph.D. degree in System Science from the University of California at Los Angeles in 1974, and has been at JPL since then. He has participated in the development of autonomous control systems flown in spacecraft for planetary missions, including the Viking mission to Mars, the Voyager mission to several planets in the outer solar system, and the Galileo mission to Jupiter and its moons. Since the mid 1980s, he has been involved in conducting research for NASA research and technology programs in robotics and control, and their application to planetary surface exploration. His research interests include analysis is hierarchical systems, estimation theory, control of nonlinear multi-body systems, and control architectures for autonomous systems. A recent area of research emphasis is that of development of system evaluation methods. His research contributions to robot dynamics led to the NASA Software of the Year Award jointly awarded to him and his JPL colleague A. Jain in 1997 (http://www.dshell.jpl.nasa.gov). (guillermo.rodriguez@jpl.nasa.gov).

Alberto Elfes obtained a Ph.D. degree in Electrical and Computer Engineering, with concentration in Robotics, from Carnegie-Mellon University in 1989. From 1986 to 1993, he held research scientist positions at the Robotics Institute and the Engineering Design Research Center, CMU, as well as at the Department of Computer Sciences, IBM T. J. Watson Research Center, where he was one of the founders of the Autonomous Robotics Laboratory. From 1994 to 2000 he was the Coordinator of Education and Research for a pilot UNDP project in strategic information technologies, and served as director of the Automation Institute, Brazil, where he established new laboratories and projects in the areas of robotics, rapid prototyping, and industrial decision support systems. During 2000/2001 he was on sabbatical in Germany as recipient of the prestigious Mercator Professorship Award. He has published extensively in the autonomous robotics area, and is the author of the Occupancy/Inference Grid framework, a stochastic approach to autonomous perception, mapping, navigation and control. In 1997 he initiated and led AURORA, the first autonomous airship project reported in the literature. He is presently a Principal Member of Technical Staff at the Jet Propulsion Laboratory, where he leads tasks for the NASA Engineering for Complex Systems (ECS) Program and the NASA Information Systems (IS) Program. His current research interests include new approaches to complex systems analysis, technology portfolio and mission design optimization; decision-theoretic approaches to risk assessment and management; and autonomy technologies for robotic UAVs, planetary aerobots, and very long distance traverse rovers. His hobbies include medieval and renaissance music, literature, and sailing. (alberto.elfes@jpl.nasa.gov)

Jeffrey H. Smith, Ph.D., is a member of the Mission and Systems Architecture Section at the Jet Propulsion Laboratory. He has developed decision analysis methodologies and tools for the analysis of advanced space missions and technologies for NASA. He has authored over 100 technical papers spanning a wide variety of disciplines and applications. Dr. Smith has been a principal investigator on computational approaches to the R&D Portfolio Selection problem and has developed algorithms and tools for technology evaluation, technology portfolio management, and automated resource allocation systems. He has extensive experience in decision support systems and software development applied to technology management. He received his BS and MS degrees in Systems Engineering from the University of Arizona; MA in Economics and Ph.D. in Business Administration from the University of Southern California, where he has taught in the Marshall School of Business and Departments of Mathematics and Systems Engineering. He has also taught at the UCLA Business and Engineering Program and the Claremont Graduate University for the Peter Drucker Executive Management Program. His hobbies include the martial arts, classical guitar, hiking, and backpacking. (jeffrey.h.smith@jpl.nasa.gov)